

Koenen, Johannes; Peitz, Martin

Working Paper

Firm Reputation and Incentives to Milk Pending Patents

CESifo Working Paper, No. 4355

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Koenen, Johannes; Peitz, Martin (2013) : Firm Reputation and Incentives to Milk Pending Patents, CESifo Working Paper, No. 4355, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/79647>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Working Papers

www.cesifo.org/wp

Firm Reputation and Incentives to “Milk” Pending Patents

Johannes Koenen
Martin Peitz

CESIFO WORKING PAPER NO. 4355
CATEGORY 11: INDUSTRIAL ORGANISATION
JULY 2013

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Firm Reputation and Incentives to “Milk” Pending Patents

Abstract

In this paper we develop a theory of patenting in which a firm preserves its reputation to only apply for a patent whenever a truly patentable idea has been generated. Firms have a short-run incentive to deviate and receive additional rents from unworthy pending patents, as well as potential rents from PTO mistakes in granting patents. We provide conditions for reputation to be preserved in equilibrium and analyze which market environments are favorable for such an equilibrium to exist. In particular, we analyze the merits of different patent systems.

JEL-Code: L150, L220, O340.

Keywords: pending patents, reputation, patent quality, patent office policy.

Johannes Koenen
Ifo Institute – Leibniz-Institute
for Economic Research
at the University of Munich
Poschingerstraße 5
Germany - 81679 Munich
koenen@ifo.de

Martin Peitz
Department of Economics
University of Mannheim
68131 Mannheim
Germany
martin.peitz@gmail.com

July 20, 2013

Martin Peitz gratefully acknowledges financial support from the German Science Foundation (SFB TR 15). He is also affiliated with CEPR, CERRE, CESifo, ENCORE, MaCCI, and ZEW.

1 Introduction

*When your competitors see the words “Patent pending” at a trade show, on your new product, on your web site, or in your sales literature, they will naturally wonder about the scope of your patent application. [...] Your patent application will not be discoverable for at least eighteen months or more, and even then, prosecution could impact what ultimately may issue. So your competitor’s fear of the unknown may provide you a temporary but substantial advantage in the marketplace. Use it well.*¹

A pending patent application is a curious thing: On the one hand, there is little to nothing that a firm’s competitors can do about it. In fact, legislation in the US even allows the details of the application to be held secret for at least 18 months, a time limit introduced only in 1999 by the American Inventors Protection Act.² On the other hand, there are a number of channels through which even the sheer evidence of a patent application may confer economic value to its holder: As in the introductory quote, competitors may change their behavior in desirable ways, such as competing less aggressively in the market or in an innovation race. Other companies may be willing to pay to purchase or license pending patents for a number of reasons: Either in order to obtain information during the period it is held secret, or later on, if for example complementary know-how of the patenting firm or institution is required for profitable use.³ Any investment that a competitor undertakes will be in the shadow of a holdup-situation arising if and when the patent is granted.⁴

In addition to reactions by competitors, a firm may benefit for other reasons. Consumers’ willingness to pay could increase if they perceive a pending patent as a signal of quality. Häussler, Harhoff, and Müller (2009) and Cockburn and MacGarvie (2009) find that a firm’s pending patents improve its access to venture capital and external fi-

¹Patent agents Robert Gunderman and John Hammond in an advice column titled “The Limited Monopoly” in the May 2007 issue of The Rochester Engineer, see www.patenteducation.com/images/200705.Limited_monopoly_-_Patent_Pending.pdf.

²See Chapter 37 §1.211 of the Code of Federal Regulations and Gallini (2002) for a brief discussion of the implications.

³Typical examples in this context are agreements with technology transfer offices of universities or government institutions. For example, the National Cancer Institute at Frederick provides “Licensing Contact Information for Patents and Patents Pending” at www-immb.ncifcrf.gov/~toms/contacts.html.

⁴See, e.g., Koenen and Peitz (2012), who also provide a broad discussion of other issues related to pending patents.

nancing.⁵ Pending patents, therefore, are valuable to firms, even if they are not granted, later on. These short-term benefits during the pending phase add to firms' temptation to apply for patents with bogus ideas. This is a serious issue, which potentially contributes to the widely documented application inflation at the Patent and Trademark Office (PTO), and the associated problems with regard to patent quality, set out, e.g., in Cook (2007) and Bessen and Meurer (2008).

This paper explores to which extent a firm's reputation may act as a countervailing force to this tendency and how patent policy can best make use of reputation mechanisms to uphold patent quality. Whether it is competitors' willingness to enter into pre-grant licensing agreements or their fear of later holdup; each source of value for pending patents requires that others believe the firm's application to be legitimate, i.e., that it will result in a patent being granted with sufficiently high probability. Intuitively, if a firm's application is rejected, its reputation will suffer so that its ability to generate rents with future pending applications diminishes or disappears.

We model this in an infinite horizon setting in the spirit of Klein and Leffler (1981) and Choi (1998): Each period, with a certain probability strictly smaller than 1, a firm generates an objectively patentable idea. Even when it does not, though, it can submit a patent application to the PTO. After some periods of inspection, the PTO grants patents to good ideas with certainty, while due to first-order mistakes it also grants patents to bad applications with a positive probability. If a patent has been granted, the firm holds it for the remainder of the patent lifetime. During the pending phase, a firm generates income from each pending patent, depending on the publicly held belief regarding patent quality, which is based on the observable history of the game. As our model focuses on inspection and PTO policy, we assume that publicly available information is limited mainly to the results of the PTO's examinations of patent applications.⁶ For each granted patent in the firm's portfolio, the firm receives a (belief-dependent) income for each period of the patent lifetime.

We study under which conditions reputational consequences, via the channel of publicly held beliefs, can induce desired behavior of firms, i.e., applying for patents only when an objectively patentable idea has been generated. This allows us to compare the

⁵Kortum and Lerner (2000) find that venture capital, in turn, may spark more innovative investment; they estimate that about 8% of innovation expenditures are linked to venture capital influx.

⁶In particular, we abstract from the information contained in potential post-grant lawsuits.

effects of different approaches to patent inspection: the “US” policy of fast inspection (small γ) with a relatively high rate of first order mistakes vs. the “European” method of more careful scrutiny at the cost of relatively slow inspection.⁷ Further, we provide a rationale for differences between values of patent applications of “young” vs. more established firms.

Related Literature To our knowledge, this is the first theoretical economic study focussing on the revenues to firms from pending patents and the ensuing incentives to apply for bogus patents. A handful of empirical studies have focussed on pending patents in an often descriptive setting, with the added difficulty that it is difficult to obtain information on unsuccessful patent applications, as Hall, Jaffe, and Trajtenberg (2005) note.

More closely related to our interest of strategic implications of pending patents, van Zeebroeck (2007) finds evidence that firms adapt their patenting strategies to benefit from the pending period of patents, using two decades worth of European application data. Firms appear to extend the pending phase, in particular by filing divisional applications which are processed more slowly, but the firm can also influence the duration of patent inspection through different ways of drafting and pursuing claims (longer specifications, number and complexity of claims, choice of route through the European Patent Office or Patent Cooperation Treaty) and by requesting (or refraining from requesting) an accelerated process. As the duration of the inspection process increases (and thereby the duration of protection through the actual patent shrinks), it becomes less likely that the firm will choose to activate the patent, once it is granted. A possible inference from this is that the pending phase may actually be *more* effective than an actual patent to protect a firm’s interests in many cases. In a similar study using US data, Popp, Juhl, and Johnson (2004) consider grant lags for US patent applications. They find that more valuable claims tend to be examined for longer periods of time by the PTO.⁸

The effects of pending patents on the performance of (mainly) young firms are studied by Häussler, Harhoff, and Müller (2009) and Cockburn and MacGarvie (2009). Both find

⁷The average time from application to decision is around two years for the US, but almost twice as long in Germany, see Hall and Harhoff (2004) who also provide empirical evidence for the difference in grant rates across the systems.

⁸This is, however, at least partially driven by the fact that patents from different areas of technology consistently differ in their inspection duration.

that pending patents lead to a higher probability of obtaining venture capital financing. According to the estimates obtained by Cockburn and MacGarvie (2009), this effect is actually stronger for pending than for granted patents.

Methodologically, our approach is related to the literature on umbrella branding, in particular Choi (1998) and Wernerfelt (1988).⁹ More broadly, our paper contributes to the literature on the firm as a bearer of reputation, which is surveyed in Bar-Isaac and Tadelis (2008). To model reputation, we follow the approach proposed by Klein and Leffler (1981) that is also used in Choi (1998).

Wernerfelt (1988) proposes a signaling model in which a firm successively produces experience goods in two periods. Even when consumers have bought the first good, they are not completely certain regarding its quality. The central result is that equilibrium beliefs exist, such that a subsequent discovery of bad quality in the second good leads to a downward re-evaluation of the quality perception of the first good, which results in revenue losses for the monopolist firm. This mechanism can lead the firm to abstain from misrepresenting the quality of its later product. As opposed to this, in Choi (1998), the quality of experience goods is revealed to consumers perfectly after purchase. The effect of a discovery of bad quality here is that consumer adjust their quality beliefs for products introduced in the future; this forward-looking effect can be efficient for the firm to refrain from extending its brand to low-quality products.

We suggest that patent applications and experience goods have a lot in common. The quality of a patent application can only be judged by competitors and the public in general after it has become accessible; for simplicity, we assume that this happens only when the PTO publishes its verdict on the application. Similar to Wernerfelt (1988), after the PTO's decision, not all uncertainty is resolved: different kinds of patent litigation such as challenges and infringement suits bear witness to the fact that even the quality of granted patents is not entirely certain. While in Choi (1998) consumers correctly assess the quality of the purchased product immediately, we introduce a delay between filing and examination as well as potential examination errors. This adds additional tradeoffs: The lower detection probability by the PTO gives firms a direct incentive to game the system; but, as in Wernerfelt (1988) there is also an indirect force at play. If the PTO makes mistakes, there is residual uncertainty regarding the quality of patents

⁹In the context of umbrella branding the firm is concerned about the reputation it enjoys with consumers, while in our model the firm is concerned about its reputation vis-a-vis its potential competitors.

even after their approval. As long as the firm has the reputation of submitting only good applications, this does not play a role, but once a firm is caught with a bad application being rejected, the beliefs regarding the existing patents will be updated, as well. The higher the likelihood of (past) mistakes by the PTO, the larger the downward correction of the beliefs with regard to patent quality, which results in a greater loss of patent value. Our model extends these approaches along a further dimension: by introducing a potential lag between application and decision (the pending period) the incentive problem of the firm is, in general, exacerbated. In the limit case in which the PTO does not commit errors, the asymmetric information problem is purely driven by uncertainty in the pending phase and reputation effects are generated only by rents accruing during the examination phase.

Our model is complementary to existing theoretical approaches that study the incentives to issue patents *without* taking the pending period into account. Horstmann, MacDonald, and Slivinski (1985) and Anton and Yao (2004) develop models in which an innovator has different ways of protecting and exploiting his intellectual property at his disposal, including trade secrecy. Integrating the pending phase into this kind of model, taking the built-in option of secrecy for an endogenous time-period into account, might shift the balance in favor of patenting. Meurer (1989), Crampes and Langinier (2002) and Koenen (2011) each study settings in which the decision to patent plays a role in signaling the quality of the underlying idea, with competitors of the firm having the chance to oppose the patent through challenges, infringement, or both. Our approach can be seen as complementary, where we focus on an earlier stage, the pending phase, and consider repeated patenting decisions.

Finally, the basic intuition that arises from our results is related to the central message of Fishman and Rob (1995). They study a competitive framework in which consumers do not observe all prices and incur search costs, while the costs of asymmetric firms vary over time. They find that larger firms which were more successful in the past are less prone to price increases after positive cost shocks. In our reputation model, we achieve a similar effect: Firms that were successful in developing ideas and patents in the past should be less tempted to endanger their reputation and play the patent system by profiting from bogus pending patents.

The plan of the paper is as follows. In Section 2 we present the model. In Section 3 we analyze the most simple case in which a patent is pending for one period, after

which its lifetime is one additional period if granted. In section 4 we analyze the general model of overlapping patents where both the patent pending phase as well as the time under patent consists of a finite number of periods. In Section 5, we apply our findings to the design of patent policy and Section 6 concludes.

2 Structure of the Model

Ideas and patent applications Consider a discrete infinite horizon setup with a single company involved in research and development. Each period t , with probability ρ the firm’s research is successful and generates a patentable idea. Regardless of success, which is private information of the firm, it can submit an application for a patent to the PTO. Therefore, each period t , the firm can take an action $a_t \in \{0, 1\}$ whether or not to submit a patent application. We denote the (publicly observable) history of actions taken by the firm until time t as A_t . In the following, we refer to the results of successful research as patentable or “good” ideas and the unsuccessful results as non-patentable or simply “bad” ideas.

Patent inspection and patent lifetime The PTO inspection technology is defined in the following way: The PTO requires γ periods to inspect a patent application, whether the underlying idea is good or bad. After this, patents are granted for good applications with certainty. Bad applications, on the other hand are not always rejected. Instead, with probability μ , the PTO commits a first-order mistake and grants a patent despite the application being bad. Detection probabilities are independent over time. If a patent is granted, the firm holds it for the remainder of the patent lifetime, i.e., $\Gamma - \gamma$ further periods.¹⁰

Publicly held beliefs of patent quality The public does not observe the true quality of an idea or patent. Instead, for each patent application, it observes the decision made by the PTO, which can be either granting a patent, rejecting an application or, if no application was submitted γ periods earlier, no decision. Each period, the public forms a belief ψ_i^t regarding the probability that an idea generated in period i was good.¹¹ This

¹⁰We discuss a micro-foundation according to which patents may be challenged or infringed below.

¹¹We omit the superscript t whenever the period in which the belief is formed is obvious from the context.

belief depends on the publicly known parameters ρ and μ as well as the observed history of the PTO decisions up to date, denoted by H_t . We discuss the formation of beliefs in equilibrium in the analysis below.

Payoffs from patents We distinguish between per-period payoffs from pending applications π and from granted patents Π . Denote the per-period rents from a good patent application during pendency in period t as $\pi_t(\psi_i; g)$, and for a bad idea as $\pi_t(\psi_i; b)$. If a patent has been granted, we denote the period t rents from a good idea as $\Pi_t(\psi_i, g)$ and a bad idea as $\Pi_t(\psi_i; b)$. If a patent application is rejected, no further rents accrue from the idea. Future payoffs are discounted at rate β .

In an imperfectly competitive market, after the patenting decision by the firm, competitors form beliefs about the quality of the patent(s) and applications, and then react according to these beliefs. We focus on the reputation mechanism and consider a reduced-form profit functions which capture market interaction between the patenting firm and its competitor.¹² We assume that, in general, each payoff increases in the public belief regarding quality, i.e., $\frac{\partial \pi(\psi_i; g)}{\partial \psi_i} \geq 0$, $\frac{\partial \pi(\psi_i; b)}{\partial \psi_i} \geq 0$, $\frac{\partial \Pi(\psi_i; g)}{\partial \psi_i} \geq 0$, and $\frac{\partial \Pi(\psi_i; b)}{\partial \psi_i} \geq 0$. Furthermore, good patents (and applications) yield a weakly higher payoff than bad patents (applications) for any given belief, i.e., $\pi(\psi_i; g) \geq \pi(\psi_i; b)$ and $\Pi(\psi_i; g) \geq \Pi(\psi_i; b) \forall \psi_i$. We note that our formulation allows reduced-form profits not to depend on the quality of the idea for given beliefs. Indeed our key insights hold even if profits given beliefs are independent of underlying patent quality. Finally, the private cost of applying for a patent is set equal to 0, as we do not want to make our arguments contingent on such costs.¹³ In the following example we provide a micro-foundation of our reduced-form profits:

Example: Consider the following micro-foundation for payoffs. Firm A, the patentee, faces potential competition from firm B, which decides whether and, if yes, when to enter the market for each product as an imitator. As a monopolist in the market, firm A receives per period profits of $\pi > 0$, which we can interpret as the willingness to pay of a unit mass of consumers net of A's costs of production. There are fixed costs of

¹²For models that explicitly take market interaction into account, see, e.g., Meurer (1989), Crampes and Langinier (2002) or Koenen (2011). Our reduced-form assumptions are compatible with a variety of models, e.g., with Cournot or Hotelling competition.

¹³However, this cost may be an interesting policy tool that may be used to deter innovations of little private value and may thus complement the reputation mechanism.

entering the market for firm B of K_e . We assume that firm B has lower marginal costs than firm A. If B enters, Bertrand competition in the product market ensues, with per period profits of $\pi_e > 0$ for B and 0 for A.

Firm A can respond to B's entry by litigation once the patent has been granted. A's lawsuit will be successful if the patent is deemed valid by the court, which results in an injunction, so that firm B must exit the market. With regard to the lawsuit, assume the following: The court will hold in A's favor, if the underlying patent is good. There are court costs of $c_B > 0$ for B. Litigation costs for A, $c_A \in \{c_H, c_L\}$ which include legal fees, opportunity costs, potential damages to the brand etc., can be either high (c_H) with probability χ , or low (c_L). χ is independent of the true quality of the patent. The realization of these costs is private information of A, while χ is common knowledge. Ex ante, B believes A's pending patent to be of high quality with probability ψ . We assume that if A's costs are high, then A doesn't litigate irrespective of whether the underlying patent is good or bad.

B can choose among the following strategies: It can enter the market during the pending phase and exit without incurring litigation once a patent is granted with the expected payoff of:

$$\sum_{i_1}^{\gamma-1} \beta^{i-1} \pi_e + (1 - \psi)(1 - \mu) \sum_{\gamma}^{\Gamma} \beta^{i-1} \pi_e - K_e. \quad (1)$$

Note that if the entry costs K_e are not too small, a critical $\bar{\psi}_1$ exists such that B prefers not entering the market if $\psi \geq \bar{\psi}_1$.¹⁴ Alternatively, B can enter and remain in the market even if the patent is granted, thereby risking litigation, with the following ex-ante expected payoff:

$$\sum_{i_1}^{\gamma-1} \beta^{i-1} \pi_e + (1 - \psi(1 - \chi)) \sum_{\gamma}^{\Gamma} \beta^{i-1} \pi_e - \beta^{\gamma} \psi(1 - \chi) c_B - K_e \quad (2)$$

Again, it is simple to derive a $\bar{\psi}_2$ equivalently to above, for which B prefers not entering the market to this strategy. Beliefs affect the competitor's choice of strategy whenever $\rho < \bar{\psi}_{1,2} < 1$. In this case, we can specify the expected profit measures introduced above as follows: $\pi_t(1; g) = \pi_t(1; b) = \Pi_t(1; g) = \Pi_t(1; b) = \pi$, $\pi_t(\rho; g) = \pi_t(\rho; b) = \Pi_t(\rho; b) = 0$

¹⁴ $\bar{\psi}_1 = [\sum_{i_1}^{\gamma-1} \beta^{i-1} \pi_e + (1 - \mu) \sum_{\gamma}^{\Gamma} \beta^{i-1} \pi_e - K_e] / [(1 - \mu) \sum_{\gamma}^{\Gamma} \beta^{i-1} \pi_e - K_e]$

and $\Pi_t(\rho; g) = (1 - \chi)(\pi - \tilde{c}_A) < \pi$, where $\tilde{c}_A$ is A 's court costs c_L discounted and annualized over the patent lifetime.

Timing The timing of each period is the following: First, the firm observes whether or not a good idea was generated this period. Then, the PTO makes the decision on the patent application submitted γ periods previously (if any). Next, the firm decides whether or not to submit an application for the current period. Finally, beliefs are updated and the payoffs for the period accrue.

3 Reputation equilibrium in the simple OLG-case

In this section we consider the simplest case, $\gamma = 1$ and $\Gamma = 2$; i.e., a patent application is pending for a single period (the application period itself), and if granted there is a single further period during which the firm obtains rents from the idea. We study the existence of a reputation equilibrium: This is a perfect Bayesian equilibrium in which the firm abstains from applying for a patent whenever a bad idea is generated. In the following, we discuss the conditions under which such an equilibrium exists.¹⁵

We start by specifying the competitors'/public's beliefs with regard to the quality of ideas. In equilibrium, i.e., as long as no deviation has been detected by an application being rejected, both applications and patents are believed to be good with certainty ($\psi = 1$). Out of equilibrium beliefs are as strict as possible: Once a patent application is rejected, competitors believe that the firm follows a "cheating" strategy and submits every idea as a patent application, irrespective of its underlying quality. For pending applications, the adjusted belief therefore is $\psi = \rho$, the publicly known probability of a good idea being generated.¹⁶ Importantly, also the perceived quality of a granted patent is adjusted downward according to Bayes' rule, due to the possibility of errors by the PTO. We denote this as $\hat{\psi} = \rho/[\mu + (1 - \mu)\rho]$.

Given that we assumed non-negative payoffs, the firm will always submit patent applications for good ideas (think of this as participation). To show that a reputation equilibrium exists and is supported by the structure of beliefs, we have to demonstrate

¹⁵Regarding the formal approach, this model is closely related to Choi (1998), who analyzes the dynamics of umbrella branding by a monopoly firm.

¹⁶The prior cannot reasonably drop below the probability of the firm generating a good idea each period, due to the assumptions regarding payoffs, i.e., $\pi(\psi_i; g) \geq \pi(\psi_i; b)$ and $\Pi(\psi_i; g) \geq \Pi(\psi_i; b) \forall \psi_i$.

that the firm prefers not to submit an application, given that it has drawn a bad idea (incentive compatibility). It is convenient to denote the expected value of the firm's revenue stream from future innovations as ν_{t+1}^g given that the patent office has made no negative decision up until period $t + 1$, and ν_{t+1}^b if an application was rejected in period $t + 1$ (or previously). This corresponds to the firm having a spotless or *good* vs. a compromised or *bad* reputation regarding its patenting decisions.

Suppose that a firm with its reputation intact has drawn a bad idea in period t . The discounted expected value of future rents $V_t(a_t, -)$ when submitting a patent application despite the underlying idea being bad is:¹⁷

$$V_t(1, -) = \pi(1; b) + \beta(\mu[\Pi(1; b) + \nu_{t+1}^g] + (1 - \mu)\nu_{t+1}^b) \quad (3)$$

The payoff from the pending patent accrues with certainty. With probability μ , the PTO fails to uncover that the patent is bad and the reputation stays intact, while with probability $1 - \mu$ the application is rejected, so that beliefs are adjusted as a response.

By contrast, if the firm refrains from submitting a bad idea as a patent in period t , the expected value of the firm is:

$$V_t(0, -) = \beta \nu_{t+1}^g \quad (4)$$

Therefore, the incentive compatibility condition $V_t(1, -) \leq V_t(0, -)$ boils down to:

$$\pi(1; b) + \beta\mu\Pi(1; b) \leq \beta(1 - \mu)[\nu_{t+1}^g - \nu_{t+1}^b] \quad (5)$$

The reputation equilibrium only exists if the short-term gains from the pending patent and, in addition, potential gains from the PTO committing an error in its evaluation are exceeded by foregone future rents due to lost reputation, weighted by the probability of being caught in the act of submitting a bad application. Explicitly, with its reputation intact and following the reputation strategy, the firm's present value of future rents is:

$$\nu_{t+1}^g = \frac{1}{1 - \beta}[\rho\pi(1; g) + \rho\beta\Pi(1; g)] \quad (6)$$

¹⁷We omit the payoff from an eventually existing patent submitted in the previous period, as it is not relevant for the decision at hand.

If, on the other hand, a bad application has been rejected, the following simple Lemma allows us to establish the discounted value of future rents:

Lemma 1: *Once an application has been rejected, it is optimal for the firm to submit an application for all ideas in the future.*

Proof: Upon rejection, beliefs are adjusted to ρ for applications and $\hat{\psi}$ for granted patents. Any further detection does not influence beliefs. Since $\pi(\rho; b) \geq 0$ and $\Pi(\hat{\psi}; b) \geq 0$, the firm chooses to submit bad applications after the first rejection. ■

Applying Lemma 1, the firm's discounted present value from future innovations given that an application was detected in the past is:

$$\nu_{t+1}^b = \frac{1}{1-\beta} [\rho(\pi(\rho; g) + \beta\Pi(\hat{\psi}; g)) + (1-\rho)(\pi(\rho; b) + \beta\mu\Pi(\hat{\psi}; b))] \quad (7)$$

Plugging into (5) gives us the following proposition:

Proposition 1: *A subgame perfect equilibrium in which firms submit only good ideas for patenting and refrain from submitting bad applications (reputation equilibrium) exists if and only if the following holds:*

$$\begin{aligned} \pi(1; b) + \beta\mu\Pi(1; b) \leq & \frac{\beta}{1-\beta}(1-\mu)[\rho\pi(1; g) - (1-\rho)\pi(\rho; b) - \rho\pi(\rho; g) \\ & + \beta(\rho\Pi(1; g) - (1-\rho)\mu\Pi(\hat{\psi}; b) - \rho\Pi(\hat{\psi}; g))] \end{aligned} \quad (8)$$

Proof: The firm always has an incentive to submit good applications. Lemma 1 establishes the optimal future choices of the firm given the original rejection. If (8) is violated, the firm will submit an application whenever it draws a bad idea, and it refrains from doing so otherwise. ■

Intuitively, the firm compares the immediate gain from the pending application (and eventually falsely granted patent) with the probability-weighted future loss from foregoing its reputation. We can observe the interesting side-effect of Lemma 1: Once a deviation has been detected, the firm has nothing left to lose and submits applications at every opportunity. This actually makes deviating more attractive, myopically, because it implies that a non-negative rent will be obtained in every future period in which a bad

idea is drawn. This effect has to be more than balanced by the expected reputational “punishment,” the loss incurred for each future good idea due to adjusted beliefs:

$$\pi(1; g) - \pi(\rho; g) + \beta\Pi(1; g) - \Pi(\hat{\psi}; g).$$

Therefore, a necessary condition for a reputation equilibrium to exist is that the probability-weighted punishment exceeds the expected gains from bad ideas once the firm’s reputation is lost, i.e.,

$$\rho\pi(1; g) - (1 - \rho)\pi(\rho; b) - \rho\pi(\rho; g) + \beta(\rho\Pi(1; g) - (1 - \rho)\mu\Pi(\hat{\psi}; b) - \rho\Pi(\hat{\psi}; g)) > 0. \quad (9)$$

With this, we obtain the following corollary:

Corollary 1: *If inequality (9) is satisfied at $\beta = 1$, there is a critical value $\bar{\beta} < 1$, such that a reputation equilibrium exists if and only if $\beta \geq \bar{\beta}$.*

Proof: Denote the left-hand side of (8) as $L(\beta)$ and the right-hand side as $R(\beta)$. If (9) holds, then $\lim_{\beta \rightarrow 1} R(\beta) \rightarrow \infty$, whereas $L(\beta)$ is bounded from above. Since $L(0) > R(0)$, by continuity of L and R in β there must exist at least one intersection and any intersection is in $(0, 1)$. Since L is linear and R convex in β , the intersection is unique and denoted by $\bar{\beta}$. ■

If the punishment effects through the reputation channel are strong enough, i.e., (9) holds at $\beta = 1$, a “patient” firm will be governed by reputational concerns and, therefore, refrain from cheating. A suitable interpretation of this parameter in the context of innovative firms is the frequency with which ideas are generated: Firms with a high innovative output and shorter intervals between ideas being generated can be interpreted as being patient—it is these firms, for which the reputation equilibrium can be supported.

Next, we turn to the inspection policy of the PTO, which in the simple setting is summarized by the probability of mistakes μ . In addition to the direct effect of a change in μ —the lower detection probability makes deviation more attractive—the difference in expected values of future ideas $[\nu_{t+1}^g - \nu_{t+1}^b]$ is affected in a less obvious way:

$$\frac{\partial[\nu_{t+1}^g - \nu_{t+1}^b]}{\partial\mu} = -\frac{\beta}{1 - \beta} \left[\rho \frac{\partial\Pi(\hat{\psi}; g)}{\partial\hat{\psi}} \frac{\partial\hat{\psi}}{\partial\mu} + (1 - \rho)(\Pi(\hat{\psi}; b) + \mu \left(\frac{\partial\Pi(\hat{\psi}; b)}{\partial\hat{\psi}} \frac{\partial\hat{\psi}}{\partial\mu} \right)) \right] \quad (10)$$

The firm's profits from future bad ideas increase, as their detection by the PTO becomes less likely; but, there are two countervailing affects which both depend on the value of granted patents. A weaker inspection regime makes the reputation mechanism more powerful; it increases the loss in value derived from future granted patents $\Pi(1, g) - \Pi(\hat{\psi}, g)$ and $\Pi(1, b) - \Pi(\hat{\psi}, b)$ after the first deviation is detected. In addition, there is a substitutional relationship between reputation and inspection quality imposed by Bayes' rule. We combine these observations into the following Lemma.

Lemma 2: *Upon rejection of a bad patent application, the value of patents granted in the future decreases. This effect is stronger for a weak patent regime (larger μ).*

Hence, from the perspective of the firm, it is worse to be caught cheating by a relatively sloppy patent office. As a thought-experiment consider two otherwise identical patenting firms, firm A in a country with a highly diligent patent office and firm B in a country with a PTO prone to mistakes. If each firm loses their reputation at the same time, we would expect the value of each of B 's future patents to drop more strongly, compared to A 's, while from then on, given the optimal strategy upon detection according to Lemma 1, the share of rejected applications should be higher for firm A than for firm B .

In the extreme, one can imagine a PTO that makes no mistakes at all, such that $\mu = 0$. Then, the critical constraint is reduced to:

$$\pi(1; b) \leq \frac{\beta}{1-\beta} [\rho\pi(1; g) - (1-\rho)\pi(\rho; b) - \rho\pi(\rho; g)] \quad (11)$$

Since all bad applications are detected, $\hat{\psi} = 1$, which shuts down reputation effects with regard to granted patents. Even then, though, a reputation equilibrium can exist. Here, the channel through which reputation takes effect is purely the difference in expected profits from pending patents.

Finally, we turn to the effects of the probability of generating a good idea ρ on the possibility to support a reputation equilibrium.

$$\begin{aligned} \frac{\partial[\nu_{t+1}^g - \nu_{t+1}^b]}{\partial \rho} = & \frac{\beta}{1-\beta} (1-\mu) [\pi(1; g) - \pi(\rho; g) + \pi(\rho; b) - \left((1-\rho) \frac{\partial \pi(\rho; b)}{\partial \rho} + \rho \frac{\partial \pi(\rho; g)}{\partial \rho} \right) \\ & + \beta \left(\Pi(1; g) - \Pi(\hat{\psi}; g) + \mu \Pi(\hat{\psi}; b) - \left((1-\rho) \frac{\partial \Pi(\hat{\psi}; b)}{\partial \hat{\psi}} \frac{\partial \hat{\psi}}{\partial \rho} + \rho \frac{\partial \Pi(\hat{\psi}; g)}{\partial \hat{\psi}} \frac{\partial \hat{\psi}}{\partial \rho} \right) \right)] \end{aligned} \quad (12)$$

We observe an indirect and a direct effect of such an increase: On the one hand, the likelihood of future punishment increases—the more likely the firm is to draw a good idea, the more sensitive it is to the expected loss of revenue. Additionally, the myopic profits from submitting bad ideas occur less often in the future. This holds both for pending and for granted patents in the future. On the other hand, the indirect effect is associated with the relative magnitude of the reputational punishment. As ρ increases, the pool of ideas from which the firm draws improves. This entails that the adjustment of beliefs in the case of detection becomes less strict and therefore the drop in value of future pending and granted patents is less steep. Applying these observations to (8), we see that for low values of ρ , the incentive constraint cannot be satisfied, because the probability of revenue losses for future good ideas is too small. On the other hand, for sufficiently high values of ρ , i.e., as $\rho \rightarrow 1$, (9) approaches 0 and the firm no longer feels the punishment effect after being exposed as a cheater. Thus, the incentive constraint can only be satisfied for intermediate values of ρ .

In practice, it may be possible to partially disentangle the indirect and direct effect. The direct effect via the likelihood of being affected by reputational punishment is related to the firm's expectations of its ideas generated in the future and may be based on information that is not necessarily publicly accessible. On the other hand, the indirect effect depends on how the public and competitors assess the firm's ability to generate good ideas—this may be related to average acceptance rates for the technological field or industry and depend less on the individual characteristics of the firm.

To briefly summarize, in the simple case with two period patents, reputational concerns can keep a firm from applying for bogus patents if the losses regarding the future income stream, weighted by the probability of detection, outweigh the immediate gains: the profits from patent pending plus the chance of a patent being granted by mistake. Being detected cheating is costly only with regard to future patentable ideas—the adjusted public beliefs result in lower rents for pending and granted patents, given that the underlying idea is truly patentable.

4 Reputation and pending patents: The general case

In this section, we consider the general case of $\Gamma > \gamma + 1 \geq 2$. This introduces two additional important effects. While in the previous setting, by construction, the firm

could have at most one active granted patent at any point in time, now it is possible for the firm to accumulate a portfolio of multiple active and pending patents. Then, the firm will have to take the effects of its decision on the value of this existing portfolio into account. The intuition is simple: In our setting, patents are comparable to experience goods. In the pending phase, their value can not be determined. Inspection grants an imperfect insight into the actual quality and, by merit of the equilibrium we are interested in, into the strategy of the firm in general. With an existing portfolio, a failure of an idea to be patented grants the public insights into the quality of the remaining pending and active patents of the firm.¹⁸

The second additional effect is related to the duration of pendency $\gamma > 1$. While in the previous section, the firm observed whether or not its cheating had been detected prior to making its next decision, in the general model it will have to decide whether to cheat again under uncertainty with regard to the outcome of the original deviation. Hence, we have to study how and whether a decision to cheat today affects the firm's decision tomorrow. We show below that if the inspection quality by the PTO is not too low, i.e., μ is not too large, then it is optimal for a firm to deviate in future periods after having deviated once, if the initial deviation has been or still can be detected. We also note that a firm that cheats today can still be rehabilitated by a combination of positive draws and errors by the PTO. We will show how this chance of rehabilitation affects the firm's incentives. In the previous section, the deviating firm needed to be "lucky" only once at the PTO to maintain the same beliefs as prior to the deviation.

To study the general case, we will require some additional notation, which we will now briefly introduce.

Beliefs in the general model As long as the reputation of the firm is intact, i.e. no application has been rejected, $\psi = 1$ for pending and granted patents, like in the simple case. Upon rejection of an application by the PTO, the public again infers that the firm's strategy is to submit applications for *every* idea. As above, this implies that all future pending patents are associated with the quality belief ρ , while the perceived quality of granted patents is $\hat{\psi} < 1$. As opposed to the case above, this shift in beliefs affects not

¹⁸This relates to work in the literature on umbrella branding, see e.g. Wernerfelt (1988) and Cabral (2000). In the previous section, a loss in reputation affects only future payoffs from ideas, in the spirit of Choi (1998).

only payoffs from future actions, but also from those in the past. Due to asymmetric information, the public cannot discern at which point in time the firm started cheating and deviated from the reputation strategy. The strictest possible inference that the public can make is that the firm has cheated from the beginning—given patent lifetimes, this means that the perceived quality of all granted patents and applications originating within Γ periods prior to the rejection of an idea will be adjusted to $\hat{\psi}$ and ρ , respectively. To spell out the firm's calculus, we will let $\Psi_s^{d,t}$ denote the beliefs that a firm expects to face from the perspective of period s , given that it has deviated in period t .

Payoffs As opposed to above, the payoffs from an idea submitted by the firm now derive from more than one period. In the case of the first deviation, reputational consequences will at the earliest ensue γ periods hence—up until then, no information relevant to belief formation will be generated from the decision to cheat. To reduce the notational burden, we introduce the following:

1. We first define the expected future payoff from a bad idea submitted in period t during pendency only (i.e., without the profits from the potentially granted patent):

$$\begin{aligned} p_t(1; b) &= \sum_{k=1}^{\gamma} \beta^{k-1} \pi(E(\psi_t^{t+k-1} | H_t, A_{t+k-1}; b)) \\ &= \sum_{k=1}^{\gamma} \beta^{k-1} \pi(1; b). \end{aligned} \tag{13}$$

Since the firm did not deviate in the past other market participants hold the belief that the pending patent is based on a good idea. This is different for ideas generated from then on, when looking more than γ periods into the future. Denote the expected cumulative probability of an application having been rejected by period s as δ_s . This object depends on detection probability, the likelihood of generating a good idea as well as the optimal strategy of the firm from the initial deviation on, which we discuss below. The expected future payoff from a pending patent application for a bad idea submitted in period $t + j$ from the perspective of period t is:

$$\begin{aligned} p_{t+j}(\Psi_t^{d,t}; b) &= \sum_{k=1+j}^{\gamma+j} \beta^{k-1} \pi(E(\psi_{t+k-1} | H_t, A_{t+k-1}; b)) \\ &= \sum_{k=1+j}^{\gamma+j} \beta^{k-1} [\delta_{t+k} \pi(\rho; b) + (1 - \delta_{t+k}) \pi(1; b)] \end{aligned} \tag{14}$$

The value depends on the expected development of the beliefs given the (expected) history of actions taken by the firm. The firm anticipates eventual future devia-

tions. The payoffs from pendency for a good idea—here, the second argument of p_{t+j} is replaced with a “ g ”—can be derived analogously.

2. Analogously, we can define the expected future payoff from a (falsely) granted patent for an idea submitted in period $t + j$ from the perspective of period t , in which the first deviation takes place as:

$$\begin{aligned} P_{t+j}(\Psi_t^{d,t}; b) &= \sum_{k=\gamma+j+1}^{\Gamma+j} \beta^{k-1} \Pi(E(\psi_{t+k-1}|A_{t+k-1}); b) \\ &= \sum_{k=\gamma+j+1}^{\Gamma+j} \beta^{k-1} [\delta_{t+k} \Pi(\hat{\psi}; b) + (1 - \delta_{t+j+k}) \Pi(1; b)] \end{aligned} \quad (15)$$

Taking the probability of PTO mistakes into account, the actual expected payoff from a granted patent of an application is $\mu P_{t+j}(\Psi_t^{d,t}; b)$.

Optimal deviation and rehabilitation The firm’s deviation today affects its choices in the future, which potentially must be taken under uncertainty with regard to the first bad application’s inspection outcome.

Proposition 2: *As long as the probability of mistakes by the PTO, μ , is not too large, then if it is optimal for a firm to submit a bad application in period t , it must also be optimal for the firm to deviate again subsequently, while the original deviation can still be detected.*

The proof is relegated to the Appendix. Intuitively, $(1 - \mu)$ is the probability that the firm has nothing left to lose due to its cheating in t , and if this is large enough, the proposition must hold.¹⁹ In the following, we will assume, that μ is not too large so that a deviation in t implies deviations in the future, given that detection is still possible.

In the simple two-period case the firm observes whether its reputation remained intact in the period after deviating, before the next action must be chosen. With probability μ , a deviation in t had no effect on the decision in $t + 1$. Proposition 2 shows, how the firm’s deviation in t affects its decisions in at least the following γ periods in the general model. Still, a return to a situation in which the original deviation no longer plays a role is possible: Suppose that a firm deviates in period t for the first time. For the next γ periods, though, it draws and submits good ideas. Then, in period $t + \gamma$, the PTO makes

¹⁹What adds an additional complication is the existence of the portfolio-effect, which we introduce in the following subsection. If detection is unlikely enough, then the positive contribution of cheating to the portfolio may turn a deviating firm into a “good citizen” concerned about its reputation.

a mistake and (falsely) grants the patent applied for in period t . In this case, the firm is indiscernible from one that has adhered to the reputation strategy throughout, and, in addition to this, it knows that no deviation will be detected based on its past actions. In this sense, a firm can be rehabilitated in our model and it takes this possibility into account when contemplating the original deviation.

The probability that a firm deviating in period t will be rehabilitated in period $t + k$ from the perspective of period t can be developed as follows: Until period $t + \gamma$, the first deviation may still be detected, therefore no sequence of draws can lead to rehabilitation. In $t + \gamma$, the firm is rehabilitated for certain, if each intermittent draw was positive and the inspection in $t + \gamma$ results in a false patent grant. The probability for this occurring from the perspective of period t is $\rho^{\gamma-1}\mu$. What happens, if there is a further bad draw in $t + 1$? The firm will deviate again, because its first deviation can still be detected. As a result, it can no longer be rehabilitated by period $t + \gamma$, but only one period later, at the earliest. Which sequence of draws and decisions by the PTO leads to this outcome? Positive draws until $t + 1 + \gamma$, as well as two false grants by the PTO: One for the application in t and one for the application in $t + 1$. The probability of this occurring from the perspective of $t + 1$ (given negative draws and deviations in t and $t + 1$) is $\mu^2\rho^{\gamma-1}$. Note that the number of paths that lead to rehabilitation by period t increases: If a firm deviating in t has been rehabilitated in period $t + \gamma$, then its decision in t no longer negatively affects its incentives to deviate in $t + \gamma + 1$. We depict the corresponding expectations of the firm by the function D_t^{t+k} : It captures the expected probability of a firm deviating in t *not* being rehabilitated in period $t + k$ from the perspective of period t . We can define this expression inductively as follows, with $X_{d,t}^{t+k} = 1 - D_t^{t+k}$.

$$X_{d,t}^{t+k} = \begin{cases} 0; & \text{if } t < \gamma \\ \rho^{\gamma-1}\mu; & \text{if } k = \gamma \\ \rho^{\gamma-1}\mu + \rho^{\gamma-1}\mu(1 - \rho)\mu; & \text{if } k = \gamma + 1 \\ X_{d,t}^{t+k-1} + \rho(X_{d,t}^{t+k-1} - X_{d,t}^{t+k-2}) + (1 - \rho)\mu(X_{d,t}^{t+k-1} - X_{d,t}^{t+k-2}); & \text{if } k \geq \gamma + 2 \end{cases}$$

Existing portfolio As discussed in the introduction to this section, the second substantial difference to the simple case is that a firm can now acquire a portfolio of pending and granted patents. Consider, for example, a granted patent, for which the firm applied in period $t - \alpha$, such that α is the age of the patent from the perspective of period t . If the

underlying idea is good, then the future revenue stream for this patent is $\Pi(\Psi_t; g)$ for the remainder of $\Gamma - \alpha$ periods. Therefore, younger ideas are more valuable than older ones and each period, as the portfolio ages, its value decreases (unless there is a new addition to it). We denote the discounted expected value of the revenue stream from pending applications and granted patents from the perspective of period t as $\Omega_t(\Psi_t, H_t, A_t)$.²⁰

The existing patent portfolio can affect the choice to deviate in t . If the firm plays the reputation strategy, then an existing patent with the remaining lifetime $\Gamma - \alpha$ will contribute a payoff of $\Pi(1; g)$ per period. What happens if the firm deviates in t and $\Gamma - \alpha > \gamma$? If the firm's cheating is detected and the public beliefs are adjusted, the payoffs from the existing patent will be adjusted downward to $\Pi(\hat{\psi}; g)$ for the remaining $\Gamma - (\alpha + \gamma)$ periods of its lifetime. For any given portfolio, $\Omega_t(1, H_t, A_t) \geq \Omega_t(\Psi_t^{d,t}, H_t, A_t)$, i.e., a deviation in t reduces the value of the existing patent portfolio in expectation. We denote this difference as $\Delta\Omega_t$. Whenever there has been a good draw in the past $\Gamma - (\gamma + 1)$ periods and if $\Pi(1; g) > \Pi(\hat{\psi}; g)$, this expression is strictly positive. A firm with an established portfolio that loses its reputation therefore can take a discrete hit to the value of its future revenues due to the portfolio effect $\Delta\Omega_t$.

The IC-constraint in the general case With these definitions in place, we can now derive the central expression that we are interested in, the IC-constraint for the general case:

Proposition 3: *A subgame perfect equilibrium in which firms submit only good ideas for patenting and refrain from submitting bad ones (reputation equilibrium) exists if and only if the following condition is satisfied for all t .*

$$\begin{aligned} & p_t(1; b) + \mu P_t(\Psi_t^{d,t}; b) + \sum_{j=1}^{\infty} D_t^{t+j} \left[\beta^j (1 - \rho) (p_{t+j}(\Psi_t^{d,t}; b) + \mu P_{t+j}(\Psi_t^{d,t}; b)) \right] \\ & \leq \sum_{j=1}^{\infty} \beta^j \rho [p_{t+j}(1; g) - p_{t+j}(\Psi_t^{d,t}; g) + P_{t+j}(1; g) - P_{t+j}(\Psi_t^{d,t}; g)] + \Delta\Omega_t \end{aligned} \quad (16)$$

Proof: Analogous to proof of Proposition 1, with Proposition 2 replacing Lemma 1. ■

²⁰For the existing patent portfolio in t , denote the number of active granted, or respectively pending, patents k periods into the future as $N_{t+k}(A_t)$ for active and $n_{t+k}(A_t)$ for pending patents. Then, for a firm on the equilibrium path, $\Omega_t(\Psi_t, H_t, A_t) = \sum_{k=1}^{\Gamma-1} \beta^k [N_{t+k} \Pi(1, g)] + \sum_{k=1}^{\gamma-1} \beta^k [n_{t+k} \pi(1, g)]$, if $\gamma > 1$; otherwise, the second term is 0. For a firm off the equilibrium path, one needs to introduce count variables to distinguish bad applications in the past.

The basic structure of (16) is similar to (8) above. As long as the firm is not rehabilitated, it reaps the profits from bad ideas (third term on the left-hand-side of the inequality). The other important change is the final term on the right-hand-side of the equation. The expected change in the value of the existing portfolio $\Delta\Omega_t$ serves to stabilize the reputation equilibrium. We will discuss how an optimal PTO policy can make use of this fact in the following section.

5 Policy Discussion

5.1 Designing the patent regime

In this section we analyze how PTO policies affect the possibility of sustaining reputation equilibria in our model. We will refer to the effects of a policy as favorable to reputation equilibria if it either decreases the left-hand-side or increases the right-hand-side of inequality (16).

Inspection Quality Much of the recent policy discussion has focussed on increasing the quality of inspection through PTOs, and comparing the inspection policies at the US and European PTOs.²¹ Adding a reputation-dimension to how we think about the patent system adds an additional facet. As examined above for the simple case, reputation can serve as a substitute for inspection through the PTO, to a certain degree (see Lemma 2). There, we differentiated between the direct effect of an increase in the likelihood of PTO mistakes μ (less likely to be caught cheating) and the countervailing indirect effect (loss of profits through reputation for each application increasing). In the general model, the direct effect is also channeled through the probability of rehabilitation. It is immediately obvious that the following must hold: $\partial D_t^{t+j}/\partial\mu \leq 0$. The probability of being rehabilitated is increasing at each point in time beyond $t + \gamma$. This channel contributes to making cheating more profitable.²² The other difference to the simple case

²¹For example, Hall and Harhoff (2004) investigate differences in grant rates between the European and the US patent offices. They show that a significant share (increasing over time up to around 30 per cent) of patent applications to the European Patent Office for ideas already successfully patented in the US were rejected. They infer that first-order mistakes play an important role for the (lack of) quality of US patents in particular.

²²A firm that is rehabilitated can generate at least weakly higher expected profits from each idea than a firm which has either lost its reputation or is under threat of losing it.

is that now a change in μ also affects the portfolio-effect $\Delta\Omega_t$. Here, again, we observe a direct and an indirect effect. In expectations, it is less likely that the PTO will discover the deviation, which diminishes the portfolio effect (direct effect). If detection takes place, though, the downward-correction in the value of the existing portfolio must be stronger for larger μ , by the familiar arguments (indirect effect). As a result, Lemma 2 immediately carries over to the general case. With respect to policy, integrating the reputation mechanism shows us that changes in inspection quality in either direction will have a smaller effect than a purely static model would predict, due to the countervailing indirect effects.

Importantly, we can make empirical predictions based on these considerations. Detection of deviations, i.e., a rejection of a patent application,²³ should affect the value of both the remaining patents owned by a firm (portfolio effect) as well as the value of future applications, through reputational contagion. This effect should be, in the sense of a difference-in-difference approach, stronger in a setting with weak patent inspection (such as the US) than in a setting with a more discerning PTO (such as the EPO). This must hold at the level of individual patents—the value of the firm is not necessarily affected in this way, due to the strategic reaction, as discussed above. Within a given patent regime, these effects should also be visible across technology sectors, if sectors can be identified, in which the quality of patent inspection is harder to uphold.

Inspection Duration A further issue regarding the improvement of inspection quality is that it comes at a cost. With given and limited resources, tougher scrutiny most likely requires examiners to spend more time with applications, which will tend to increase the average time required for inspection, which in our model is represented by γ . There are a number of important issues related to the duration of patent pendency, for an overview we refer to Koenen and Peitz (2012). To be able to cleanly separate the effects, we proceed as follows: We first analyze the case of a shift from γ to $\gamma' = \gamma + 1$ with perfect inspection ($\mu = 0$), and then demonstrate the additional tradeoffs arising from dropping this restriction.

As a first step, let us consider the case with $\mu = 0$. Here, rehabilitation is impossible and the development of beliefs $\Psi_t^{d,t}$ is not stochastic but completely predetermined, as the firm loses its reputation when the first bad application is examined in period γ/γ'

²³Or a comparable event, such as a court declaring a patent void.

with certainty.

Corollary 2: *For $\mu = 0$, an increase in the patent inspection duration γ destabilizes the reputational equilibrium iff the following condition holds:*

$$\beta^{\gamma+1}\pi(1; b) + \sum_{j=0}^{\gamma+1} \beta^{\gamma+1} [\rho(\pi(1; b) - \pi(\rho; b)) + (1 - \rho)(\pi(1; g) - \pi(\rho; g))] + \sum_{j=1}^{\infty} \beta^{\gamma+1+j} [(1 - \rho)\pi(\rho; b) + \rho(\pi(\rho; g) - \Pi(1; g))] > 0 \quad (17)$$

Proof: Rewrite (16) as $C(\gamma) > 0$ and $C(\gamma') > 0$, respectively, and take the difference. ■

Note that due to $\mu = 0$, $\Pi(1; g) = \Pi(\rho; g)$, i.e., this case also allows us to isolate the effects of and on pending patents, the main focus of this study. This also implies that no portfolio effect exists. Analyzing (17), we can differentiate two effects. There is a type of short-term lump-sum benefit at $\gamma + 1$, which is embodied in the first two terms: The firm receives the benefit of the original bad application for an additional period (instead of having it rejected, earlier); more than this, it loses its reputation one period later, and therefore receives higher benefits on all other pending applications filed in the meantime for one additional period. The third term resembles the continued difference in expected payoffs over time: Each future bad patent application will be discovered one period later, and therefore provide additional rents. On the other hand, for future good applications, the firm suffers from lost reputation for one additional period—in the alternate regime, the patent would already have been granted and there would no longer have been a reputational effect.

An immediate side-effect of this adjustment in γ is that the relative importance of the granted patent with regard to payoffs diminishes as opposed to the profits from pendency. Since patent applications generate value during the pending phase, we should see this reflected in firms' IP strategies (an issue we abstracted from in our formal analysis), as pending durations are increasing over time in general, and are substantially higher in Europe than in the US, in particular.²⁴ Note that each of these effects holds for the case of $\mu > 0$, as well.

²⁴Evidence which is consistent with such strategic behavior is provided by van Zeebroeck (2007): A substantial share of *granted* European patent applications, whose average pendency in the sample exceeds five years, do not become patents since applicants do not pay the required fees. While compatible with strategic behavior, other explanations can be given. In particular, firms holding pending patents may wait for good news about the private value of their idea and enter the patenting process only after receiving favorable news: The longer they wait the less likely that good news arrives.

For $\mu > 0$, the portfolio effect decreases in γ , due to aging. One final effect is introduced due to the fact that the earliest potential detection takes place one period later.²⁵ The direct result of this, which we observed above, is that reputational punishment is deferred, which makes deviation more attractive. However, an increase in γ makes rehabilitation less likely, because more consecutive positive draws are necessary for this to occur.²⁶

Pre-Grant Reviews and Intensity of Competition Throughout our entire analysis, it is the reputational “punishments” ($p(1, g) - p(\rho; g)$ and $P(1, g) - P(\rho; g)$), which support the existence of the reputation equilibrium. These differences depend on how strongly profits react to the beliefs held by the public. The more aggressive the reaction of competitors to a deterioration of said beliefs, the more careful the patenting firm will be to protect its reputation. In general, therefore, competitiveness of the industry in question will support the reputation equilibrium. Further, due to the central role that the myopic benefits from pending bad applications play in destabilizing the reputation equilibrium, policies that aim at reducing $\pi(., b)$ are especially beneficial.

Corollary 3: *More intense competition as well as policies that reduce the gains from pending patents for bad ideas, $\pi(., b)$ are favorable to the existence of reputation equilibria.*

In this regard, recent reform steps towards introducing pre-grant reviews initiated by third parties are an important step in such a direction: Without this instrument, pending patents are virtually unassailable. Critics might note that it is not only bad ideas that will be attacked, thereby raising the costs of patent applications which might hamper innovation. Note that from a reputational perspective, this might even be conducive, as long as firms with an intact reputation are less likely to have their applications challenged: In this case, the pre-grant review process increases the difference $\pi(1; g) - \pi(\Psi^{d,t}; g)$. Further note that such a wedge should be larger for sectors in which competition is fiercer, as there are more strategic players interested in launching reviews.

²⁵All else given, therefore, beliefs with regard to patent quality are weakly higher under the slower regime.

²⁶The latter effect intuitively derives from the fact that the option-value of returning to the reputation strategy must be at least weakly greater than zero. If the probability of being able to execute this option decreases, then this affects the profitability of deviating adversely.

5.2 Sophisticated policies toward patentees

Due to legal restrictions, there is very limited scope in practice as to discriminating between applicants in the level of scrutiny exerted by the PTO. However, it is instructive to initially abstract from legal restrictions and to allow for discriminatory policies, before discussing sophisticated policies which do not explicitly discriminate between different types of firms. In particular, we discuss an immediate implication of the portfolio effect, namely that a patent office with limited resources should treat firms implicitly or explicitly differently depending on their patent portfolio.

Suppose first that a PTO can only condition on the age of the firm. As we have seen above, an older firm has an in expectation larger patent portfolio than a young firm. This implies, that a non-discriminatory PTO policy, i.e., one that checks all patent applications with the same rigor, leads to the incentive constraint of old firms to be often slack. Hence, a resource-constrained PTO designs its optimal discriminatory policy such that young firms' patent applications are, on average, investigated more thoroughly than older firms.

While discriminating based on firm age constitutes an improvement over a non-discriminatory policy, the PTO can use its resources even more efficiently by conditioning its investigation of patent applications on the existing patent portfolio. It can avoid that an old firm with a series of bad draws facing a lax investigation of its applications deviates from the proposed equilibrium path, applies with all new ideas and, thus, destroys its reputation. A resource-restricted PTO designs its optimal discriminatory policy such that a firm of a given age is investigated less thoroughly if it has applied more often for a patent over the last Γ periods and if it holds a large stock of granted patents. Taking Corollary 2 and the discussion below it into account, we can further formulate a simple decision rule for PTO examiners: If two applications arrive simultaneously, but have to be examined sequentially, the application of the firm, which has the smaller stock of patents and applications should be examined first.

Does one observe such discriminatory inspection behavior in practice? While there is, as of yet, no evidence regarding such behavior by PTOs, a number of empirical studies points towards discriminatory inspection behavior playing an important role in the context of venture capital decisions and firm valuations. If the value of a firm is composed to a large degree of patent applications, then the role of a venture capitalist or bank deciding whether or not to finance the firm is very similar to the role of the

PTO. Our analysis above would let us expect substantial non-linearities of the firm value with respect to granted and pending patents depending on the grant history. In particular, early granted patents may carry additional value due to their contribution to the firm's reputation. This exact pattern is detected by Greenberg (2009) in a study of the valuations of Israeli high-tech companies. We would further expect that for younger firms, i.e., firms without an established reputation, there should be a substantial difference in the value of pending vs. granted patents, if only the latter contribute to reputation. Greenberg (2009) find that the effect of granted patents is about twice as large as that of pending patents for young firms. As opposed to this, we would expect this difference to shrink substantially once the firm has established its reputation, and, in fact, almost 90% of the difference disappears for older firms. This overall pattern suggests that firm reputation plays an important role, as recognized by investors and the market. Other studies to discover substantial non-linearities in the values of patents of firms at early stages include Cockburn and MacGarvie (2009) and Häussler, Harhoff, and Müller (2009). However, none of these studies specifically focusses on the reputation effects of patents. We would like to see empirical research in this regard.

Finally, our analysis provides a different motivation for why two-tiered systems of patent validity might be effective (see, e.g., the discussion in Lichtman and Lemley (2007)): In our model, firms whose (lack of) portfolio fails to introduce slack to their IC-constraint would opt into a stricter, more intense and in this sense more expensive examination. On the other hands, firms owning a more extensive portfolio would not require this additional signal.

6 Conclusion

Pending patents generate significant economic value to applicants, as has been empirically demonstrated by papers such as Greenberg (2009), Cockburn and MacGarvie (2009), Häussler, Harhoff, and Müller (2009) or Gans, Hsu, and Stern (2008): They boost the valuations of startups, increase the likelihood of receiving venture capital, and pave the way for licensing agreements. Firms can unilaterally generate them by filing applications, upon which they remain secret and non-opposable for a significant amount of time (at least 18 months). While this may appear cynical in the face of the overload-problem suffered by PTOs, as discussed by Caillaud and Duchene (2011), this raises the

provocative question: Why don't we observe *more* abuse of pending patents?

In this paper, we develop a simple theory focusing on the benefits from pending patents to a firm. We identify reputation as a countervailing effect to the inflation of dubious patent applications. A firm that expects to be applying for patents in the future may want to refrain from abusing the system today if this leads to negative inferences by market participants in the case of future detection by the PTO. Reputation then deters firms from milking the ability to create pending patents. Sustaining reputation equilibria is desirable from a policy perspective.

In this paper we show the existence of reputation equilibria—i.e., we only consider equilibria in which the reputation of a firm remains intact, as the incentive compatibility constraint is satisfied throughout from the beginning and, therefore, for any potential path. This serves to stress our main point, but may be criticized for its lack of realism. In particular, in our model the portfolio effect only plays a role in formulating sophisticated PTO policies.

A straightforward extension would be to let firms be “born” in period 0 with existing portfolios. In this case, firms whose portfolio falls short of a certain threshold might not be able to commit to the reputation strategy.

To accommodate a richer set of effects without such an ad hoc assumption on existing portfolios, future work may want to consider a stochastic extension in which firms sacrifice their reputation if they draw a sufficiently bad productivity or demand shock. Firms then do not regain their reputation. New firms with sufficiently good realizations of their demand and/or productivity parameters enter the patent system only with good ideas and, thus, obtain a positive reputation. In this scenario, due to the portfolio effect, even with somewhat worse realizations in the future, the IC-constraint will still be satisfied. However, if the realization is sufficiently bad or the firm suffered a string of bad luck in the generation of ideas, the IC constraint may fail to hold. This may make reputation eventually unsustainable along the equilibrium path and lead to firm turnover such that firms exploit their reputation when receiving a sufficiently bad realization and, then, exit the market.

Finally, we would welcome empirical analyses of the relationship between reputation and the value of pending and granted patents. Reputation is a potential explanation for observed non-linearities in valuations. Differences across patent regimes, technological sectors and competitiveness between industries, as discussed above, may provide the

required variation to explore these issues in greater detail.

A Appendix

A.1 Proof: Optimal deviation

In the following, we derive conditions under which a firm submitting a bad idea in period t (deviating) chooses to deviate in the future. We will call a firm that returns to the reputation strategy (i.e., which submits only patentable ideas) “rehabilitated”. Conceptually, rehabilitation will necessarily be possible: Consider a firm which has submitted a faulty idea in period t , but experiences only good draws in the following γ periods, which it submits to the PTO. Then, in period $t + \gamma$, the faulty idea is examined. If the period- t deviation is not detected, then the firm does not have to fear detection of any past delinquent behavior and can be rehabilitated and return to the reputation strategy. Given observable history, the firm cannot be distinguished from one that has not deviated (and, indeed, has added $\gamma - 1$ good applications to its portfolio).

Therefore, we proceed as follows: We first derive a condition that ensures that if a firm deviated in t , it will again deviate in $t + 1$, given a bad draw in this period. We then extend this condition recursively to period $t + n$, with $n \leq \gamma$, and also take potential positive draws in the meantime into account. Suppose that the firm deviated in period t . Then, in period t , the following inequality must hold:

$$\begin{aligned}
 & p_t(1; b) + \mu P_{t+1}(\Psi_t^{d,t}; b) + \sum_{j=1}^{\infty} D_t^{t+j} \left[\beta^j (1 - \rho) (p_{t+j}(\Psi_t^{d,t}; b) + \mu P_{t+j}(\Psi_t^{d,t}; b)) \right] \\
 & - \sum_{j=1}^{\infty} \beta^j \rho_{t+j} [p_{t+j}(1; g) - p_{t+j}(\Psi_t^{d,t}; g) + P_{t+j}(1; g) - P_{t+j}(\Psi_t^{d,t}; g)] - \Delta \Omega_t \geq 0.
 \end{aligned} \tag{18}$$

Consider the incentives that the firm faces in the next period. We define the objects $\Omega'_t(-)$ and $\Omega'_t(+)$ as the value of the patent portfolio Ω_t in terms of its discounted expected revenue stream, after aging each item within by one period and adding one new bad $(-)$ or good $(+)$ item to it in period t , *while assuming that the deviation will not be detected by the PTO*—the importance of this assumption will become clear shortly: In principle, adding a bad idea to the patent portfolio has two potential incentive effects: If the deviation is *not* detected, then this additional item contributes to the portfolio effect and may enhance the incentives to comply with the reputation strategy. If the deviation *is* detected, then the firm’s reputation breaks down and it has no further incentive to abstain from submitting any idea.

To address the adjusted portfolio effect, consider the case that the deviation in t

will *not* be detected γ periods hence. Suppose that the firm draws a bad idea in $t + 1$, again. For this case we derive the new incentive-compatibility constraint. Note that all elements are identical to equation 18, due to the assumption of non-detection, except for the portfolio effect. Now consider the difference between the IC-constraints in t and $t + 1$, subtracting the latter from the former. The difference between the two, which we refer to as $G_{t,t+1}$ is:

$$G_{t,t+1} = \Delta\Omega_{t+\gamma} - \Delta\Omega'(-)_{t+\gamma} \quad (19)$$

$G_{t,t+1}$ denotes the change in the portfolio-effect between periods t and $t + 1$. This object typically will be negative, as can be illustrated with the following examples. First take the case in which $p^- = p^+$ and $P^- = P^+$, i.e., given non-detection, objectively non-patentable ideas are as valuable as good ones. If a firm has a full portfolio at t —i.e., has submitted a patent in every period in the relevant past—then, given non-detection and the assumptions on p and P , the change in the portfolio effects, $G_{t,t+1}$, must be 0: The firm has a full portfolio at the start of both periods, with one idea dropping out, each idea aging by one period and the new (bad) idea entering.²⁷ The aging of the existing portfolio then exactly counteracts the entry of the new idea regarding the portfolio effect. For non-full portfolios, though, generally the addition to the portfolio will outweigh the portfolio-aging effect. In the “worst case”, there was no portfolio in t and only the one bad idea has been added, so that the portfolio aging does not countervail the addition of the bad idea at all. Therefore, $G_{t,t+1}$ must be bounded from below by $-\beta^\gamma [P(1; b) - P(\rho; b) - (\Pi(1; b) - \Pi(\rho; b))/\beta]$.

Now consider the case in which the deviation in t *will* be detected γ periods later. In this case, reputation breaks down in $t + \gamma$, independent of the firm’s decision in $t + 1$. Deviation in $t + 1$ has no reputational consequences, and therefore the firm will deviate again if the expected revenue stream from the idea is positive, i.e., if the following holds:

$$\sum_{j=1}^{\gamma-2} \beta^{j-1} \pi(1; b) + \beta^{\gamma-1} \pi(\rho; b) + \mu \sum_{k=1}^{\Gamma-\gamma} \beta^{\gamma+k-1} \Pi(\rho; b) \geq 0, \quad (20)$$

which we denote as $B_{t,t+1} > 0$. By our assumptions on payoffs, the lefthand-side of (20)

²⁷One can immediately see that for other assumptions on the structure of payoffs, in particular if $P^-(1) - P^-(\hat{\psi}) < P^+(1) - P^+(\hat{\psi})$, the change in the portfolio effect can also be negative. The arguments below imply that in these cases deviation in t implies deviation in $t + 1$ for any μ .

must be strictly positive, even for $\mu \rightarrow 0$. The condition for a deviation in t implying deviation in $t + 1$ therefore can be expressed as:

$$\mu G_{t,t+1} + (1 - \mu) B_{t,t+1} \geq 0 \tag{21}$$

As $B_{t,t+1}$ is strictly positive and $G_{t,t+1}$ is bounded from below, this inequality must hold as long as μ is not too large. To complete the argument, we can now adjust (19) for any intermittent draw: A positive draw increases the difference in the portfolio effect more strongly than a negative draw, but nevertheless the difference will have to be bounded from below. Therefore, a sufficiently low μ will ensure that a deviation in the past which can still be detected (cf. the argument in the first paragraph above) implies that the firm will deviate again at the next opportunity.

References

- ANTON, J. J., AND D. A. YAO (2004): “Little Patents and Big Secrets: Managing Intellectual Property,” *RAND Journal of Economics*, 35, 1–22.
- BAR-ISAAC, H., AND S. TADELIS (2008): “Seller Reputation,” *Foundations and Trends in Microeconomics*, 4, 273–351.
- BESSEN, J. E., AND M. J. MEURER (2008): *Patent Failure*. Princeton University Press.
- CABRAL, L. (2000): “Stretching Firm and Brand Reputation,” *RAND Journal of Economics*, 31, 658–673.
- CAILLAUD, B., AND A. DUCHENE (2011): “Patent Office in Innovation Policy: Nobody’s Perfect,” *International Journal of Industrial Organization*, 29(2), 242–252.
- CHOI, J. P. (1998): “Brand Extension as Informational Leverage,” *Review of Economic Studies*, 65, 655–669.
- COCKBURN, I. M., AND M. J. MACGARVIE (2009): “Patents, Thickets and the Financing of Early-Stage Firms: Evidence from the Software Industry,” *Journal of Economics & Management Strategy*, 18, 729–773.
- COOK, J. P. (2007): “On Understanding the Increase in U.S. Patent Litigation,” *American Law and Economics Review*, 9(1), 48–71.
- CRAMPES, C., AND C. LANGINIER (2002): “Litigation and Settlement in Patent Infringement Cases,” *RAND Journal of Economics*, 33(2), 258–74.
- FISHMAN, A., AND R. ROB (1995): “The Durability of Information, Market Efficiency and the Size of Firms,” *International Economic Review*, 36, 19–36.
- GALLINI, N. T. (2002): “The Economics of Patents: Lessons from Recent U.S. Patent Reform,” *Journal of Economic Perspectives*, 16, 131–154.
- GANS, J., D. HSU, AND S. STERN (2008): “The Impact of Uncertain Intellectual Property Rights on the Market for Ideas: Evidence for Patent Grant Delays,” *Management Science*, 54, 982–997.

- GREENBERG, G. (2009): “Small Firms, Big Patents? Estimating Patent Value Using Data on Start-Ups’ Financing Rounds,” *mimeo*.
- HALL, B. H., AND D. HARHOFF (2004): “Post-Grant Reviews in the U.S. Patent System - Design Choices and Expected Impact,” *Berkeley Technology Law Journal*, 19(3), 989–1015.
- HALL, B. H., A. JAFFE, AND M. TRAJTENBERG (2005): “Market Value and Patent Citations,” *RAND Journal of Economics*, 36, 16–38.
- HÄUSSLER, C., D. HARHOFF, AND E. MÜLLER (2009): “To Be Financed or Not... - The Role of Patents for Venture Capital Financing,” *ZEW - Centre for European Economic Research Discussion Paper*, 09-003(7115).
- HORSTMANN, I., G. M. MACDONALD, AND A. SLIVINSKI (1985): “Patents as Information Transfer Mechanisms: To Patent or (Maybe) Not to Patent,” *Journal of Political Economy*, 93, 837–858.
- KLEIN, B., AND K. B. LEFFLER (1981): “The Role of Market Forces in Assuring Contractual Performance,” *Journal of Political Economy*, 89, 615–641.
- KOENEN, J. (2011): “Patent Litigation and Patent Reforms,” *mimeo*.
- KOENEN, J., AND M. PEITZ (2012): “Firm Reputation and the Incentives to Milk Pending Patents,” in *Recent Advances in the Analysis of Competition Policy and Regulation*, ed. by J. E. Harrington, and Y. Katsoulacos, chap. 3, pp. 49–74. Edward Elgar.
- KORTUM, S., AND J. LERNER (2000): “Assessing the Contribution of Venture Capital to Innovation,” *RAND Journal of Economics*, 31, 674–692.
- LICHTMAN, D., AND M. A. LEMLEY (2007): “Rethinking Patent Law’s Presumption of Validity,” *Stanford Law Review*, 60(1), 45–72.
- MEURER, M. J. (1989): “The Settlement of Patent Litigation,” *RAND Journal of Economics*, 20(1), 77–91.

- POPP, D., T. JUHL, AND D. K. JOHNSON (2004): “Time in Purgatory: Examining the Grant Lag for U.S. Patent Applications,” *Topics in Economic Analysis & Policy*, 4.
- VAN ZEEBROECK, N. (2007): “Patents Live Only Twice: A Patent Survival Analysis of the Determinants of Examination Lags, Grant Decisions and Renewals,” *CEB Working Paper 07/028*.
- WERNERFELT, B. (1988): “Umbrella Branding as a Signal of New Product Quality: An Example of Signalling by Posting a Bond,” *RAND Journal of Economics*, 19, 458–466.